arXiv:1409.2175v2 [math.OC] 16 Oct 2014

No-Free-Lunch Theorems in the Continuum

Aureli Alabert
Department of Mathematics
Universitat Autònoma de Barcelona
08193 Bellaterra, Catalonia
Aureli.Alabert@uab.cat

Alessandro Berti
Dipartimento di Matematica
Università di Padova
35121 Padova, Italy
berti@math.unipd.it

Ricard Caballero
Department of Mathematics
Universitat Autònoma de Barcelona
08193 Bellaterra, Catalonia
rcaballero@mat.uab.cat

Marco Ferrante
Dipartimento di Matematica
Università di Padova
35121 Padova, Italy
ferrante@math.unipd.it

October 17, 2014

Abstract

No-Free-Lunch Theorems state, roughly speaking, that the performance of all search algorithms is the same when averaged over all possible objective functions. This fact was precisely formulated for the first time in a now famous paper by Wolpert and Macready, and then subsequently refined and extended by several authors, always in the context of a set of functions with discrete domain and codomain. Recently, Auger and Teytaud have shown that for continuum domains there is typically no No-Free-Lunch theorems. In this paper we provide another approach, which is simpler, requires less assumptions, relates the discrete and continuum cases, and that we believe that clarifies the role of the cardinality and structure of the domain.

Keywords: No-Free-Lunch, stochastic processes, black-box optimisation.
Mathematics Subject Classification (2010): 68Q25 (60G, 90C26)

1 Introduction

In [7], Wolpert and Macready formulated rigorously a principle which was already intuitively known to the operations research practitioners: All search or optimization algorithms perform equally well when their performance is averaged against all possible objective functions. This principle has been known since then as the *No-Free-Lunch Theorem* (NFL for short).

The precise formulation of the Wolpert-Macready NFL Theorem will be stated in Section 2 (Theorem 2.1), but the basic assumptions are that we are dealing with the set of all functions $f\colon \mathcal{X} \to \mathcal{Y}$ between two finite sets $\mathcal{X}$ and $\mathcal{Y}$, and that the "averaging" is uniform over all these functions. The measure of performance can be any function of the images $f(x_1), \ldots, f(x_m)$ of the points $x_1, \ldots, x_m$ sampled by the algorithm.

In [5], Schumacher, Vose and Whitley extended the result to some subsets of all functions (those called "closed under permutation"), whereas Igel and Toussaint [4] stated it for some non-uniform measures. The language of probability theory allows to formulate these statements in a unified and easier way and, as we will see, it is absolutely necessary to switch from the discrete (finite) setting to the continuum. In [2], Auger and Teytaud considered for the first time this case, and their result is essentially negative: No NFL theorems exist in the continuum.

Our goal in this paper is to improve and clarify the results of Auger and Teytaud, particularly Theorem 4.1, [2]. First of all, we show situations where NFL theorem do exist. This apparent paradox is resolved by noticing that the hypotheses imposed in [2] invalidate our examples. In fact, the authors seem to specifically look for conditions under which no NFL theorem can hold true. The theorem is indeed correct, although there is a gap in the proof, as explained in Section 3. We must also point out, however, that their paper contains much more material of interest on this and other matters.

The point of view adopted here is different: We establish a simple and natural definition of the NFL property and look for the necessary conditions implied by this definition. In this sense, our main result is Theorem 3.9. The conclusion we reach is that there are no No-Free-Lunch theorems for functions whose domain is the real number system, except for a few extreme cases. Admittedly, the theorem contains an additional mild technical condition (the existence of second-order moments) that we have not yet been able to remove.

The relevance of this theoretical discussion for the field of global *black-box* optimisation comes from the so-called *probabilistic models*: In many practical optimisation problems there is little information about the objective function, with no access to derivatives or to any explicit formula; we are only allowed to ask the function for its value at a point of our choice and, after observing the value returned, we may decide on the next point to sample the function; and so on. Moreover, function evaluations can be expensive, and we are constrained to make only a small number of them. In these cases, it may be useful to think that the function has been drawn at random from some set of functions, according to some probability law (perhaps with some unknown parameters) that one specifies using prior information. Technically, we are then in the presence of a *stochastic process*, from which our function is a particular path. Different algorithms will choose different points for the successive evaluations, and some may perform better than others by exploiting better the model, *unless* there is a No-Free-Lunch theorem for that model. If this is the case, all algorithms perform the same and, in particular, pure blind search is as good as any other proposal. In the present paper we will see that the presence of the No-Free-Lunch property reduces to a few probabilistic models, which are not really important in practice.

The paper is organised as follows: In Section 2 we state the definitions and preliminaries both from algorithmics and from probability theory that are strictly needed in the rest of the paper. In Section 3 we state the main results: We show that No-Free-Lunch cases do exist in the continuum; we impose then a hypothesis of measurability of the stochastic process involved, and we see that NFL can only appear if we are essentially in a discrete setting (Theorem 3.5), or the model consists of a trivial constant process (Theorem 3.9). In Section 4 we justify the investigation of the existence (or not) of NFL properties in the continuum and propose some open questions.

2 Preliminaries

We follow approximately the notations of [4] and [2], with some convenient modifications.

2.1 Algorithmic concepts

Let $\mathcal{X}$ and $\mathcal{Y}$ be any two sets. The set of all functions $f\colon \mathcal{X} \to \mathcal{Y}$ can be identified with the Cartesian product $\mathcal{Y}^{\mathcal{X}}$. Denote

$$\mathcal{E}_0 := \{\emptyset\},\ \mathcal{E}_1 := \mathcal{X} \times \mathcal{Y},\ \ldots,\ \mathcal{E}_m = \mathcal{X}^m \times \mathcal{Y}^m \tag{1}$$

and $\mathcal{E} := \bigcup_{m \geq 0} \mathcal{E}_m$.

A *(random) algorithm* A is a mapping $A\colon \mathcal{Y}^{\mathcal{X}} \times \mathcal{E} \times \Theta \to \mathcal{E}$, where $(\Theta, \mathcal{G}, Q)$ is a probability space and, if $e = ((x_1, y_1), \dots, (x_m, y_m)) \in \mathcal{E}_m$, then $A(f, e, \theta) \in \mathcal{E}_{m+1}$ and

$$A(f, e, \theta) = ((x_1, y_1), \dots, (x_m, y_m), (x_{m+1}, y_{m+1})) \tag{2}$$

with $y_{m+1} = f(x_{m+1})$. Therefore, we can think of $A(f, e) := (X_1, Y_1), \dots, (X_m, Y_m))$ as a random vector $\Theta \to \mathcal{E}$. This definition formalises the fact that the algorithm chooses the next point based on the previous points and an (optional) random mechanism represented by the probability space $(\Theta, \mathcal{G}, Q)$. One may assume that f is never evaluated more than once at the same point.

A *measure of performance* of the algorithm is any function C of the values obtained by evaluating f during the algorithm. Formally,

$$C\colon \bigcup_{m\geq 1} \mathcal{Y}^m \to \mathbb{R}\ . \tag{3}$$

A typical measure of performance for optimization problems is the function $C(y_1, \dots, y_m) = \min\{y_1, \dots, y_m\}$, the best observed value after m evaluations. (Notice that the measures of performance we are talking about are not related to algorithmic complexity, e.g. to the number of evaluations needed to reach the end of a procedure.)

To state the basic Wolpert-Macready Theorem, rephrased in our probability-theoretic language, consider another probability space $(\Omega, \mathcal{F}, P)$, and a random variable $f\colon \Omega \to \mathcal{Y}^{\mathcal{X}}$. Now f is random and $f(\omega)$, for each ω, is a specific function $\mathcal{X} \to \mathcal{Y}$. Denote by $A^m\colon \mathcal{Y}^{\mathcal{X}} \times \Theta \to \mathcal{E}_m$ the successive application, m times, of algorithm A to the initial empty sequence $\emptyset$, and by $A^m_Y\colon \mathcal{Y}^{\mathcal{X}} \times \Theta \to \mathcal{Y}^m$ its second component. Finally, let us abbreviate $y := (y_1, \dots, y_m) \in \mathcal{Y}^m$.

The conditional probability $Q\{A^m_Y(h, \theta) = y \;\big/\; f(\omega) = h\}$ is the probability that the algorithm A produce the particular sequence of function values $y := (y_1, \dots, y_m)$ when applied to the function h. This probability is either 0 or 1 for deterministic algorithms.

Theorem 2.1. (Wolpert-Macready [7], Theorem 1).
Assume that $\mathcal{X}$ and $\mathcal{Y}$ are finite sets. Let $f\colon \Omega \to \mathcal{Y}^{\mathcal{X}}$ be a random variable that chooses functions $f(\omega) \in \mathcal{Y}^{\mathcal{X}}$ with the uniform discrete probability law. That means, for every $h \in \mathcal{Y}^{\mathcal{X}}$,

$$P\{\omega \in \Omega :\ f(\omega) = h\} = |\mathcal{Y}|^{-|\mathcal{X}|}\ ,$$

where $|\cdot|$ denotes cardinality.

Then, the law of A^m_Y is the same for all algorithms. Precisely stated: let A and B be two algorithms; then, for all $m \in \mathbb{N}$ and all $y \in \mathcal{Y}^m$,

$$[P \times Q]\{(\omega, \theta) :\ A^m_Y(f(\omega), \theta) = y\} = [P \times Q]\{(\omega, \theta) :\ B^m_Y(f(\omega), \theta) = y\}\ . \tag{4}$$

Since the law of f is uniform,

$$\begin{aligned}
&[P \times Q]\{(\omega, \theta) :\ A^m_Y(f(\omega), \theta) = y\} \\
&= \sum_{h \in \mathcal{Y}^{\mathcal{X}}} Q\{\theta :\ A^m_Y(h, \theta) = y \;\big/\; f(\omega) = h\} \cdot P\{\omega : f(\omega) = h\} \\
&= |\mathcal{Y}|^{-|\mathcal{X}|} \sum_{h \in \mathcal{Y}^{\mathcal{X}}} Q\{\theta :\ A^m_Y(h, \theta) = y \;\big/\; f(\omega) = h\}\ .
\end{aligned}$$

Therefore, under the hypotheses of the theorem, (4) can be written

$$\sum_{h \in \mathcal{Y}^{\mathcal{X}}} Q\{\theta :\ A^m_Y(h, \theta) = y \;\big/\; f(\omega) = h\} = \sum_{h \in \mathcal{Y}^{\mathcal{X}}} Q\{\theta :\ B^m_Y(h, \theta) = y \;\big/\; f(\omega) = h\} \tag{5}$$

for any two algorithms A and B, and for all $y \in \mathcal{Y}^m$, $m \in \mathbb{N}$.

And still another, informal, way to formulate the result of Wolpert and Macready is that if an algorithm performs better than pure blind random search in a particular set of functions, then it must perform worse than random search *in the mean* on the complementary set.

For simplicity, we will assume that we deal with deterministic algorithms from now on, although everything can be easily extended to accommodate random algorithms. For deterministic algorithms, equality (5) can be written

$$\left|\{h \in \mathcal{Y}^{\mathcal{X}} : \ A_Y^m(h) = y\}\right| = \left|\{h \in \mathcal{Y}^{\mathcal{X}} : \ B_Y^m(h) = y\}\right| .$$

2.2 Probability concepts

For the sake of completeness and the reader's convenience, we summarise here, albeit in a very compact way, all concepts from measure and probability theory that are used in the sequel.

Let $(\Omega, \mathcal{F}, P)$ be a probability space and $(S, \mathcal{S})$ any measurable space. An S-valued random variable f is a measurable mapping $f\colon \Omega \to S$. A *random function* is simply a random variable with values in the space of all functions between two sets $\mathcal{X}$ and $\mathcal{Y}$, the second one equipped with some σ-field; that is, we take $S = \mathcal{Y}^{\mathcal{X}}$, and the natural choice for $\mathcal{S}$ is the *product σ-field*, i.e. the smallest σ-field that turns every projection $\mathcal{Y}^{\mathcal{X}} \to \mathcal{Y}$ into a measurable mapping. Random functions are also called *stochastic processes*, especially when $\mathcal{X}$ is an interval I of the real line and $\mathcal{Y}$ is the set of real numbers $\mathbb{R}$ endowed with the Borel σ-field. We assume in the rest of the paper that $\mathcal{X} = [0,1]$ and $\mathcal{Y} = \mathbb{R}$, so that we are dealing with random functions $f(\omega)\colon [0,1] \to \mathbb{R}$.

The *law* of a stochastic process is the image measure of P through the mapping f. That means, it is the probability μ on $(S, \mathcal{S})$ such that

$$\mu(B) = P(f^{-1}(B)) \ , \quad \forall B \in \mathcal{S} \ .$$

Given a stochastic process f, the composition $f(t) := \delta_t \circ f$ of f with the Dirac delta at $t \in [0,1]$ is automatically a real random variable $f(t)\colon \Omega \to \mathbb{R}$, with respect to the Borel σ-field on $\mathbb{R}$. The random vectors of the form $(f(t_1), \dots, f(t_m))$, with $t_1, \dots, t_m \in [0,1]$, are the finite-dimensional projections of f. Their laws, the *finite-dimensional distributions*, determine the law of the whole process.

A stochastic process can be represented in different ways: As a function-valued random variable, as above, or as a family of real-valued random variables, $\{f(t),\ t \in [0,1]\}$, or as a mapping from the product $\Omega \times [0,1]$ into $\mathbb{R}$, defined in the obvious way: $(\omega, t) \mapsto f(t, \omega)$. A process is said to be *measurable* if, in the last representation, it is a measurable mapping $\Omega \times [0,1] \to \mathbb{R}$ when $\Omega \times [0,1]$ is endowed with the product σ-field $\mathcal{F} \times \mathcal{B}([0,1])$, where $\mathcal{B}([0,1])$ denotes the Borel σ-field of $[0,1]$. Stochastic processes mentioned in Auger–Teytaud [2] are always considered measurable. This is an important hypothesis in their results, and its role will be made clear in the present paper. A process is said to be of *second-order* if all its variables are square-integrable, implying that they have finite expectation and variance.

With some abuse of notation, we use the same symbol f to denote several different related objects: f is a function $\Omega \to \mathbb{R}^{[0,1]}$, or $\Omega \times [0,1] \to \mathbb{R}$; for every $\omega \in \Omega$, $f(\omega)$ is a function $[0,1] \to \mathbb{R}$; for all $t \in [0,1]$, $f(t)$ is a random variable $\Omega \to \mathbb{R}$; and finally, for all t and ω, the value $f(t, \omega)$ is a real number.

We will also use occasionally the customary abbreviations a.s. for *almost surely* (i.e. true with probability 1), and a.e. for *almost everywhere* (i.e. true except a set of measure zero with respect to Lebesgue measure).

3 Main results

Recall, from the notations in Section 2, that $A_Y^m(f(\omega))$ is the random vector consisting of the images $(f(t_1),\dots,f(t_m))(\omega) \in \mathbb{R}^m$ produced by applying m iterations of algorithm A on the function $f(\omega) \in \mathbb{R}^{[0,1]}$.

Definition 3.1. *Let C be a performance measure, measurable on $\mathbb{R}^m$, for all m. We say that a stochastic process $f = \{f(t),\ t \in [0,1]\}$ satisfies the* No-Free-Lunch property with respect to *C if for any two algorithms A and B, and for all $m \in \mathbb{N}$, the random variables*

$$\omega \mapsto C\big(A_Y^m(f(\omega))\big) \quad \textit{and} \quad \omega \mapsto C\big(B_Y^m(f(\omega))\big)$$

have the same law.

Intuitively, the NFL property states that the information about C that we get after having sampled m points is the same no matter which algorithm we use. In particular, blind search performs as well as any other algorithm.

Definition 3.2. *We say that a stochastic process $f = \{f(t),\ t \in [0,1]\}$ satisfies the* No-Free-Lunch property *if it does so with respect to all possible performance measures C.*

Remark 3.3. It is easily seen (see e.g. Auger and Teytaud [2], Lemma 2.3), that if f satisfies the No-Free-Lunch property then the random variables

$$\omega \mapsto A_Y^m(f(\omega)) \quad \text{and} \quad \omega \mapsto B_Y^m(f(\omega))$$

have the same law. Conversely, if the above random variables have the same law, then f satisfies the No-Free-Lunch property for all performance measures C. One may say that NFL is a extremely strong form of stationarity. A *stationary process* has invariant laws under translations: The law of $(f(t_1),\dots,f(t_n))$ and $(f(t_1+h),\dots,f(t_m+h))$ are the same, for every h and any dimension m, provided all indices belong to the set where the process is defined, the interval $[0,1]$ in our case. It is clear that the NFL property is much stronger.

Examples 3.4. *We can readily show two examples of random functions enjoying the No-Free-Lunch property:*

1. *Consider a set of indices I of arbitrary cardinality, and a family of random variables $f_i\colon \Omega \to \mathbb{R}$, $i \in I$, mutually independent and identically distributed, with any non-degenerate probability law. (As a specific case, consider for instance Bernoulli or Gaussian variables, and $I = \mathbb{R}^+$ as the set of indices).*

   Such a stochastic process exists by the classical Kolmogorov Extension Theorem (see e.g. [1]). All n-dimensional joint distributions are the direct product measure of the individual laws, and are therefore the same. The NFL property is then trivially satisfied.

2. *Consider any random variable X and define a constant process $f(t) \equiv X$, for all $t \in [0,1]$. The NFL property is also immediate to check.*

These examples show that there exist NFL situations also in the continuum case, and that the cardinality of I alone cannot be the responsible of the lack of No-Free-Lunch.

It is certainly true that a continuous-time stochastic process with all variables mutually independent can hardly be of any interest in modelling a real phenomenon (in sharp contrast with the discrete case). Notice for example, that in the common case of Gaussian variables, almost all sample paths (i.e., all functions in the set we are considering, with probability 1) are

unbounded from below and from above, which makes pointless to search for or to approximate the minimum value. As another example, if the variables have the uniform law in an interval $[\alpha, \beta]$, then almost all sample paths are bounded, with infimum equal to α and supremum equal to β, although the probability that these values are attained is zero. A very different problem is the case when α or β are unknown and we try to estimate them by sampling m points from independent variables distributed uniformly in $[\alpha, \beta]$; this is indeed a statistical problem of a real practical interest.

Such trivial NFL situations do not appear if we impose on f the condition of being a measurable process. This is the main result of Auger and Teytaud [2]. We reformulate it as the problem of finding a necessary condition for having NFL in a measurable process, and show that in this case we are dealing essentially with a constant process. This possibility does not appear in [2], because of the hypothesis of existence of a so-called "proper median", that the authors introduce in the definition of NFL, and that it looks somewhat artificial. We will not use this concept. We also point out that the argument in [2] is in our opinion not complete, since at some point in the proof of their Theorem 4.1 there is a confusion between the underlying randomness of the process and the eventual randomness of the algorithm applied.

We start by showing that measurability and independence together collapses the process to an essentially discrete-time stochastic process.

Theorem 3.5. *Let $f = \{f(t),\ t \in [0,1]\}$ be a measurable stochastic process with mutually independent random variables $f(t)$. Then, for almost all t with respect to Lebesgue measure, the random variable $f(t)$ is constant with probability 1.*

Proof. We treat first the particular case in which the process is bounded and centred. Then the result will be easily extended to the general case.

First case: Assume $\mathrm{E}[f(t)] = 0$ and that there is a constant $m \in \mathbb{N}$ such that for all t, $|f(t)| < m$.

If $f\colon \Omega \times [0,1] \to \mathbb{R}$ is a measurable mapping, then the partial mappings $f(\omega)\colon [0,1] \to \mathbb{R}$ are also measurable, for $\omega \in \Omega$ almost surely. Since, moreover, all the sample paths are bounded, it makes sense to consider their Lebesgue integrals

$$g(t) = \int_0^t f(s)\, ds\ , \quad t \in [0,1]\ .$$

From the properties of the integral, $g(t)$ is a process with continuous paths almost surely. We may leave it undefined for the exceptional set of probability zero, because this is not relevant.

The random variables $g(t)$ are also clearly bounded (e.g. by m itself), and therefore the second moment $\mathrm{E}[g(t)^2]$ is finite. But we see that it is in fact equal to zero:

$$\mathrm{E}[g(t)^2] = \mathrm{E}\Big[\int_0^t f(s)\, ds \cdot \int_0^t f(r)\, dr\Big] = \mathrm{E}\Big[\int_0^t \int_0^t f(s)f(r)\, dsdr\Big] = \int_0^t \int_0^t \mathrm{E}[f(s)f(r)]\, dsdr =$$
$$= \int_0^t \int_0^t \mathrm{E}[f(s)^2] \cdot \mathbf{1}_{\{s=r\}}\, dsdr + \int_0^t \int_0^t \mathrm{E}[f(s)]\,\mathrm{E}[f(r)] \cdot \mathbf{1}_{\{s\neq r\}}\, dsdr\ ,$$

where the interchange of integral and expectation is justified by the boundedness of all functions involved, which allows to apply the Fubini theorem, and we have used the hypothesis $\mathrm{E}[f(s) \cdot f(r)] = \mathrm{E}[f(s)] \cdot \mathrm{E}[f(r)]$. Now, the first integral is equal to zero because we are integrating over the line $\{s = r\}$, which has zero Lebesgue measure, and the second one is also zero because the variables are centred.

The equality $\mathrm{E}[g(t)^2] = 0$ implies that for all $\omega \in \Omega$ except maybe in a subset $N_t \subset \Omega$ of probability zero, one has $g(t, \omega) = 0$. In particular, this is true for all $t \in \mathbb{Q}$ and, since $\mathbb{Q}$ is a countable set, we have that $N = \bigcup_{t\in\mathbb{Q}} N_t$ has probability zero. By the continuity of the paths,

we obtain that for $\omega \in \Omega - N$, $g(t,\omega) = 0$ for all $t \in [0,1]$. The integral being zero for all t, the integrand is also zero except maybe on a set of Lebesgue measure equal to zero. We conclude that a.s., and for almost all $t \in [0,1]$, $f(t,\omega) = 0$.

General case:
Let $m \in \mathbb{N}$, and define:

$$\bar{f}_m(t) := f(t) \cdot \mathbf{1}_{\{|f(t)|<m\}}$$

and

$$\hat{f}_m(t) = \bar{f}_m(t) - \mathrm{E}[\bar{f}_m(t)] \ .$$

The random variables $\hat{f}(t)$ are also mutually independent. Applying the particular case above we have that for all m, the law of $\hat{f}_m(t)$ is a Dirac delta at zero, for $t \in [0,1]$ a.e., and therefore $\bar{f}_m(t)$ is constant ω-a.s, t-a.e.

Now, let $N_{t,m} \subset \Omega$ be the set where $\bar{f}_m(t) \neq f(t)$. The random variable $f(t)$ is almost surely equal to a constant $a(t,m)$ on the set $\Omega - N_{t,m}$, which tends, as $m \to \infty$, to a set $\Omega - N_\infty(t)$, whose probability is 1, given that $\lim_{m\to\infty} P\{|f(t)| \leq m\} = 1$. We obtain that the constant $a(t,m)$ cannot depend on m, and conclude that for almost all $t \in [0,1]$, with respect to Lebesgue measure, the random variables $f(t)$ are degenerated. □

In other words, Theorem 3.5 states that, under the hypotheses of measurability and independence, randomness can only appear on a time set of zero Lebesgue measure. Hence, the process can be considered, in a measure-theoretic sense, as a discrete time stochastic process.

For the remaining of the section, we assume that we deal with second-order processes. We will show, in Theorem 3.9 below, that a measurable, second-order process satisfying the NFL property is trivial: All their random variables are almost surely equal.

First, we state some preliminary results in the form of lemmas:

Lemma 3.6. *Let $(\Omega, \mathcal{F}, P)$ be a probability space and $f : \Omega \times [0,1] \to \mathbb{R}$ a measurable second-order stochastic process. Then the NFL property can be satisfied only if the random variables $f(t)$ are identically distributed and the covariance $\mathrm{Cov}(f(t), f(s))$ is constant for all t, $s \in [0,1]$, $t \neq s$.*

Proof. If the NFL property is satisfied, then by Remark 3.3 the vectors $A_Y^m(f)$ and $B_Y^m(f)$ are identically distributed for any $m \in \mathbb{N}$ and any pair of algorithms A and B.

Take $m = 1$. Given two values t and s in $[0,1]$, let A be a deterministic algorithm that chooses t as initial point and B another deterministic algorithm that chooses s as initial point. Then $f(t)$ and $f(s)$ are identically distributed.

Take now $m = 2$. Given two couples of different points (t_1, t_2) and (s_1, s_2), let A be a deterministic algorithm that choses (t_1, t_2) as the first two points and B be a deterministic algorithm that chooses (s_1, s_2). Then the random vectors $(f(t_1), f(t_2))$ and $(f(s_1), f(s_2))$ are identically distributed. This fact implies in particular that

$$\mathrm{Cov}(f(t_1), f(t_2)) = \mathrm{Cov}(f(s_1), f(s_2)) \ .$$

□

Lemma 3.7. *Let $X_1, \ldots, X_n$ ($n \geq 2$) be real square-integrable random variables with a common covariance $\rho = \mathrm{Cov}[X_i, X_j]$ when $i \neq j$, and define $V := \max_i \mathrm{Var}[X_i]$, the maximum of their variances. Then*

$$\rho \geq -\frac{V}{n-1}$$

Proof. We have

$$0 \le \mathrm{Var}\Big[\sum_{i=1}^{n} X_i\Big] = \sum_{i=1}^{n} \mathrm{Var}[X_i] + \sum_{\substack{i,j=1 \\ i\neq j}}^{n} \mathrm{Cov}[X_i, X_j] \le Vn + n(n-1)\rho$$

and the result follows at once. □

The proof of the next result is immediate:

Lemma 3.8. *Let $f\colon \Omega \times [0,1] \to \mathbb{R}$ be a second-order stochastic process such that the variables $f(t)$ are identically distributed, with finite common mean μ and positive common variance σ^2. Then f satisfies the NFL property if and only if the same holds for $(f-\mu)/\sigma$.*

By Lemmas 3.6, 3.7 and 3.8, we can restrict the search of a second-order measurable stochastic process satisfying the NFL property to the case when the variables $f(t)$ are identically distributed, with zero mean, unit variance and such that for some $\rho \ge 0$, $\mathrm{Cov}[f(t), f(s)] = \rho$ for any pair $t \neq s$.

In the next theorem we use Fourier analysis, following quite closely some arguments that can be found in Crum [3], to prove our main result:

Theorem 3.9. *Let $\{f(t,\omega) : \ t \in [0,1], \ \omega \in \Omega\}$ be a measurable, second-order stochastic process, with $\mathrm{E}[f(t)] = 0$ and $\mathrm{E}[f(t)^2] = 1$ for all $t \in [0,1]$, satisfying the NFL property, and defined in some probability space $(\Omega, \mathcal{F}, P)$.*

Then, the process is constant, in the sense that there exists a random variable $X\colon \Omega \to \mathbb{R}$ such that $P\{\omega \in \Omega : \ f(t,\omega) = X(\omega)\} = 1$, $\forall t \in [0,1]$.

Proof. We are going first to extend the process from $[0,1]$ to the whole real line, in order to apply Fourier transform techniques comfortably. Define:

$$f(k+t) := \begin{cases} f(t), \ \text{if } t \in (0,1], \ k = 1, 2, \dots \\ f(t), \ \text{if } t \in [0,1), \ k = -1, -2, \dots \end{cases}$$

From the previous lemmas, we know that the covariance function $K(t) := \mathrm{E}[f(t+s)f(s)]$ of the extended process is equal to some ρ ($0 \le \rho \le 1$) for t in an interval $(-\varepsilon, \varepsilon)$, and all $s \in \mathbb{R}$, except for $t = 0$, in which is equal to 1. Our purpose is to see that in fact ρ must be equal to 1, from which the conclusion will be easily drawn.

It can be readily seen that the extended process is also measurable, using that a set $A \in \mathcal{B}(I) \otimes \mathcal{F}$, where I is any interval in $\mathbb{R}$, is also $\mathcal{B}(\mathbb{R}) \otimes \mathcal{F}$-measurable when considered as a subset of $\mathbb{R} \times \Omega$. This implies that, for almost all ω, $t \mapsto f(t,\omega)$ is a Borel measurable function, that it makes sense to consider the integrals

$$\int_{\mathbb{R}} e^{-2s^2} f(s)^2 \, ds \ ,$$

and that they are measurable functions $\Omega \to \mathbb{R}$. Taking expectation and applying Fubini's theorem,

$$\mathrm{E}\Big[\int_{\mathbb{R}} e^{-2s^2} f(s)^2 \, ds\Big] = \int_{\mathbb{R}} e^{-2s^2} \, \mathrm{E}[f(s)^2] \, ds = \int_{\mathbb{R}} e^{-2s^2} \, ds < \infty \ .$$

This means that, ω-a.s., $s \mapsto e^{-s^2} f(s)$ belongs to $L^2(\mathbb{R})$. It also belongs to $L^1(\mathbb{R})$, ω-a.s:

$$\mathrm{E}\Big[\int_{\mathbb{R}} e^{-s^2} |f(s)| \, ds\Big] = \int_{\mathbb{R}} e^{-s^2} \, \mathrm{E}[|f(s)|] \, ds \le \int_{\mathbb{R}} e^{-s^2} \, \mathrm{E}[f(s)^2]^{1/2} \, ds = \int_{\mathbb{R}} e^{-s^2} \, ds < \infty \ .$$

For such ω, consider the function

$$g(s) := e^{-s^2} f(s) - e^{-(s+t)^2} f(s+t) \ .$$

Since $g \in L^1(\mathbb{R})$, we may take its Fourier transform

$$\hat{g}(\xi) = \int_{\mathbb{R}} g(s) \cdot e^{-2\pi i s\xi} \, ds$$

which can be written as

$$\int_{\mathbb{R}} e^{-s^2} f(s) \cdot e^{-2\pi i s\xi} \, ds - \int_{\mathbb{R}} e^{-s^2} f(s) \cdot e^{-2\pi i (s-t)\xi} \, ds \ ,$$

or

$$(1 - e^{-2\pi i t\xi}) \cdot \hat{F}(\xi) \ ,$$

where $\hat{F}$ is the Fourier transform of $s \mapsto e^{-s^2} f(s)$.

Since $g \in L^2(\mathbb{R})$ also, by Plancherel's Theorem,

$$G(t,\omega) := \int_{\mathbb{R}} g(s)^2 \, ds = \int_{\mathbb{R}} |\hat{g}(\xi)|^2 \, d\xi = \int_{\mathbb{R}} |1 - e^{-2\pi i t\xi}|^2 \cdot |\hat{F}(\xi)|^2 \, d\xi \ .$$

The integrand tends to zero as $t \to 0$, and it is dominated by $4 \cdot |\hat{F}(\xi)|^2 \in L^1(\mathbb{R})$. Therefore, $\lim_{t\to 0} G(t,\omega) = 0$, a.s.

Moreover,

$$\begin{aligned} G(t,\omega) &= \int_{\mathbb{R}} \left| e^{-s^2} f(s) - e^{-(s+t)^2} f(s+t) \right|^2 ds \\ &\leq 2\Big[\int_{\mathbb{R}} e^{-2s^2} f(s)^2 \, ds + \int_{\mathbb{R}} e^{-2(s+t)^2} f(s+t)^2 \, ds \Big] = 4 \int_{\mathbb{R}} e^{-2s^2} f(s)^2 \, ds \ , \end{aligned}$$

that belongs to $L^1(\Omega)$, as we have seen before. By the Dominated Convergence Theorem again,

$$\lim_{t\to 0} \mathrm{E}[G(t,\omega)] = 0 \ .$$

On the other hand,

$$\begin{aligned} G(t,\omega) &= \int_{\mathbb{R}} e^{-2s^2} f(s)^2 \, ds + \int_{\mathbb{R}} e^{-2(s+t)^2} f(s+t)^2 \, ds - 2 \int_{\mathbb{R}} e^{-s^2-(s+t)^2} f(s) f(s+t) \, ds \\ &= 2 \int_{\mathbb{R}} e^{-2s^2} f(s)^2 \, ds - 2 \int_{\mathbb{R}} e^{-s^2-(s+t)^2} f(s) f(s+t) \, ds \ . \end{aligned}$$

The expectation of the first term is equal to

$$2 \int_{\mathbb{R}} e^{-2s^2} \, ds = \sqrt{2\pi} \ .$$

For the second, it yields

$$2 \int_{\mathbb{R}} e^{-s^2-(s+t)^2} K(t) \, ds = \sqrt{2\pi} e^{-t^2/2} K(t)$$

(the interchange of integral and expectation is justified here by checking first the integrability).

We get

$$\mathrm{E}[G(t,\omega)] = \sqrt{2\pi}(1 - e^{-t^2/2} K(t)) \ .$$

Hence

$$0 = \lim_{t\to 0} \mathrm{E}[G(t,\omega)] = \lim_{t\to 0} \sqrt{2\pi}(1 - e^{-t^2/2} K(t)) \ ,$$

which implies that $\lim_{t\to 0} K(t) = 1$, and we conclude that $\rho = 1$, as we wanted to see.

Finally, since $\mathrm{Cov}[f(t), f(s)] = 1$, we have $f(t) = \alpha f(s) + \beta$, for some α, β. But the variables are centred, and this implies $\beta = 0$, whereas the unit variances yield $|\alpha| = 1$. The negative value of α is impossible because the covariance is nonnegative. Hence $f(t) = f(s)$, almost surely, for all t and s, and the proof is complete. □

Corollary 3.10. *In view of Lemma 3.8, the conclusion of Theorem 3.9 is true without the hypotheses of null expectation and unit variance.*

Notice that the conclusion of the previous theorem and corollary does not mean that almost all sample paths are constant, because the null set N_t where the equality $f(t) = X$ fails depends on t. However, in an optimisation setting it is natural to specify a regularity assumption on the functions, besides the probabilistic model. For example, the continuity of the paths (or simply the right or left continuity) automatically yields that the union $\cup_{t\in[0,1]}N_t$ has probability zero, and in that case one may say that the process is constant in the sense that, except on a set $N \subset \Omega$ of probability zero, all paths $f(t)$ are constant.

Summarizing the present section:

- Without measurability assumptions, we showed two examples of NFL property in the continuum: The case in which all variables are independent, identically distributed, and the case where the process is constant: $f(t) = X$ a.s, $\forall t$, for some random variable X. Both are unimportant from the optimisation practitioner's point of view, and both are ruled out in the mentioned paper by Auger and Teytaud, by imposing the measurability and the "proper median" hypotheses, respectively.
- We have shown that measurability and independence together lead to a "discrete time" process, and that measurability and NFL (for second-order processes) imply that the process is constant. With the three conditions together, or simply measurability, independence and stationarity, one gets easily that each variable of the process must be almost surely equal to some constant k, the same for all of them: $P\{f(t) = k\} = 1$, $\forall t$.

4 Conclusion and open questions

It is frequently argued that in realistic scenarios the hypotheses of the NFL theorems are always violated, already in the genuinely discrete cases. We have shown that in the continuum the necessary conditions for NFL are even stronger and far too restrictive to be found in practice.

This means that in every practical situation there must be some information on the objective function that permits, in principle, to choose algorithms that perform better than pure blind search. We believe that the usefulness of the (no)-NFL statements is precisely on the theoretical side, to highlight that any proposal of a search algorithm, supported by a benchmark of functions in which it behaves well, should be accompanied by a study of the benchmark common features that help that algorithm beat the others. In other words, as has been emphatically pointed out in a recent expository article [6]: «It is clear now that for the practitioner the correct question is not *which algorithm I have to use* but first of all *what is the geometry of the objective function*».

Before paper [2], there was, to our knowledge, no special interest in investigating the existence of NFL theorems in the continuum. The typical argument was that only the discrete case is important in practice, since the computations are always made in finite precision. This is no longer true, since arbitrary precision computations are commonplace today, meaning that the underlying computational model is at least the rational number system $\mathbb{Q}$.

There are still two questions that deserve further study concerning NFL theorems, both in the discrete and in the continuum settings:

- The first one is the consideration of noisy functions, that means, black-box functions that may answer differently when asked twice for the value at the same point t. This is not uncommon in practice, since the computation of the objective value at a feasible point may

involve itself some randomness or the heuristic solution of another optimisation problem. In that case, the algorithms to consider should be allowed to sample more than once the same point.

- More importantly, the second question refers to the concept of No-Free-Lunch itself. As we have seen, the NFL property is so strong that constant processes are the only measurable processes that qualify. But if we are just concerned with minimizing a function, the relevant performance measure is $C(y_1, \ldots, y_n) = \min\{y_1, \ldots, y_n\}$, or perhaps some related function. Recalling Definition 3.1 applied to this measure, it is easy to see that NFL with respect to C is not sufficient to conclude that the laws of $\omega \mapsto A_Y^m(f(\omega))$ and $\omega \mapsto B_Y^m(f(\omega))$ have to be the same, and then our Theorem 3.9 need not be true. We believe that this point deserves further investigation.

5 Acknowledgements

This work has been supported by grants numbers MTM2011-29064-C03-01 from the Ministry of Economy and Competitiveness of Spain; UNAB10-4E-378, co-funded by the European Regional Development Fund (ERDF); and 60A01-8451 from the University of Padova.

References

[1] Robert B. Ash. *Real analysis and probability.* Academic Press, New York-London, 1972. Probability and Mathematical Statistics, No. 11.

[2] Anne Auger and Olivier Teytaud. Continuous lunches are free plus the design of optimal optimization algorithms. *Algorithmica*, 57(1):121–146, 2010.

[3] M. M. Crum. On positive-definite functions. *Proc. London Math. Soc. (3)*, 6:548–560, 1956.

[4] Christian Igel and Marc Toussaint. A no-free-lunch theorem for non-uniform distributions of target functions. *J. Math. Model. Algorithms*, 3(4):313–322, 2004.

[5] C. Schumacher, M. D. Vose, and L. D. Whitley. The no free lunch and problem description length. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2001*, pages 565–570. Morgan Kaufmann, 2001.

[6] Loris Serafino. Derivatives: What does the no free lunch theorem actually say? *Notices of the AMS*, 61(7):750–755, 2014.

[7] David H. Wolpert and William G. Macready. No free lunch theorems for optimization. *IEEE TRANSACTIONS ON EVOLUTIONARY COMPUTATION*, 1(1):67–82, 1997.